arXiv:1608.06072v2 [cs.LG] 3 Oct 2016

Uniform Generalization, Concentration, and Adaptive Learning

Ibrahim Alabdulmohsin
CEMSE Division, KAUST Thuwal 23955, Saudi Arabia

Abstract

One fundamental goal in any learning algorithm is to mitigate its risk for overfitting. Mathematically, this requires that the learning algorithm enjoys a small generalization risk, which is defined either in expectation or in probability. Both types of generalization are commonly used in the literature. For instance, generalization in expectation has been used to analyze algorithms, such as ridge regression and SGD, whereas generalization in probability is used in the VC theory, among others. Recently, a third notion of generalization has been studied, called *uniform generalization*, which requires that the generalization risk vanishes uniformly in expectation across all bounded parametric losses. It has been shown that uniform generalization is, in fact, *equivalent* to an information-theoretic stability constraint, and that it recovers classical results in learning theory. It is achievable under various settings, such as sample compression schemes, finite hypothesis spaces, finite domains, and differential privacy. However, the relationship between uniform generalization and concentration remained unknown. In this paper, we answer this question by proving that, while a generalization in expectation does not imply a generalization in probability, a *uniform* generalization in expectation does imply concentration. We establish a chain rule for the uniform generalization risk of the composition of hypotheses and use it to derive a large deviation bound. Finally, we prove that the bound is tight.

1 INTRODUCTION

One of the central questions in statistical learning theory is to establish the conditions for generalization from a finite collection of observations to the future. Mathematically, this is formalized by bounding the difference between the empirical and the true risks of a given learning algorithm $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \to \mathcal{H}$, where $\mathcal{Z}$ is the observation space and $\mathcal{H}$ is the hypothesis space.

Informally, suppose we have a learning algorithm $\mathcal{L}$ that receives a sample $S_m = \{Z_1, \ldots, Z_m\}$, which comprises of m i.i.d. observations $Z_i \sim p(z)$, and uses S_m to select a hypothesis $H \in \mathcal{H}$. Because H is selected *based on* the sample S_m, its *empirical* risk on S_m is a biased estimator of its *true* risk with respect to the distribution of observations $p(z)$. The difference between the two risks, referred to as the *generalization* risk, determines the prospect of over-fitting in the learning algorithm.

In the literature, generalization bounds are often expressed either in expectation or in probability. Let $L(\cdot; H) : \mathcal{Z} \to [0, 1]$ be some parametric loss function that satisfies the Markov chain $S_m \to H \to L(\cdot; H)$. We write $R_{true}(H)$ and $R_{emp}(H; S_m)$ to denote, respectively, the true and the empirical risks of the hypothesis H w.r.t. $L(\cdot; H)$:

$$R_{emp}(H; S_m) = \frac{1}{m} \sum_{Z_i \in S_m} L(Z_i;\, H)$$
$$R_{true}(H) = \mathbb{E}_{Z \sim p(z)} \big[L(Z;\, H) \big] \qquad (1)$$

Then, generalization in expectation and generalization in probability are defined as follows:

Definition 1 (Generalization in Expectation)**.** *The expected generalization risk of a learning algorithm $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \to \mathcal{H}$ with respect to a parametric loss $L(\cdot; H) : \mathcal{Z} \to [0, 1]$ is defined by:*

$$R_{gen}(\mathcal{L}) = \mathbb{E}_{S_m, H|S_m} \big[R_{emp}(H; S_m) - R_{true}(H) \big], \quad (2)$$

where the empirical risk $R_{emp}(H; S_m)$ and the true risk $R_{true}(H)$ are given by Eq. (1). A learning algorithm $\mathcal{L}$ generalizes in expectation if $R_{gen}(\mathcal{L}) \to 0$ as $m \to \infty$ for all distributions $p(z)$.

Definition 2 (Generalization in Probability)**.** *A learning algorithm $\mathcal{L}$ generalizes in probability if for any positive constant $\epsilon > 0$, we have:*

$$p\Big\{\big|R_{true}(H) - R_{emp}(H; S_m)\big| > \epsilon\Big\} \to 0 \quad \textit{as } m \to \infty,$$

where the probability is evaluated over the random choice of S_m and the internal randomness of $\mathcal{L}$.

Clearly, for bounded loss functions, a generalization in probability implies a generalization in expectation but the converse is not generally true.

In general, both types of generalization have been studied in the literature. For instance, generalization in probability is used in the Vapnik-Chervonenkis (VC) theory, the covering numbers, and the PAC-Bayesian framework, among others (Vapnik, 1999; Blumer et al., 1989; McAllester, 2003; Bousquet and Elisseeff, 2002; Bartlett and Mendelson, 2002; Bousquet et al., 2004; Audibert and Bousquet, 2007). Generalization in expectation, on the other hand, was used to analyze learning algorithms, such as the stochastic gradient descent (SGD), differential privacy, and ridge regression (Hardt et al., 2016; Dwork et al., 2015; Shalev-Shwartz and Ben-David, 2014). Its common tool is a replace-one averaging lemma, similar to the Luntz-Brailovsky theorem (Luntz and Brailovsky, 1969; Vapnik and Chapelle, 2000), which relates generalization to algorithmic stability (Hardt et al., 2016; Shalev-Shwartz and Ben-David, 2014). Generalization in expectation is often simpler to analyze, but it provides a weaker performance guarantee.

Recently, however, a third notion of generalization has been introduced in Alabdulmohsin (2015), which is called *uniform* generalization. It also expresses generalization bounds in expectation, but it is stronger than the notion of generalization in Definition 1 because it requires that the generalization risk vanishes uniformly across all bounded parametric loss functions. Importantly, uniform generalization is shown to be *equivalent* to an information-theoretic algorithmic stability constraint, and that it recovers classical results in learning theory. It has been connected to the VC dimension as well (Alabdulmohsin, 2015). Moreover, many conditions can be shown to be sufficient for uniform generalization. These include differential privacy, sample compression schemes, perfect generalization, robust generalization, typical generalization, finite description lengths, or finite domains. Indeed, we prove in Appendix A that all such conditions are sufficient for uniform generalization.

Unfortunately, uniform generalization bounds hold only in expectation without any concentration guarantees. This sheds doubt on the utility of the notion of uniform generalization and its information-theoretic approach of analyzing learning algorithms. For instance, we will later construct a learning algorithm that generalizes *perfectly* in expectation w.r.t. a *specific* parametric loss even though it does not generalize almost surely over the random draw of the sample S_m. Hence, generalization in expectation is insufficient to ensure that a generalization will take place in practice.

Nevertheless, we will establish in this paper that a *uniform* generalization in expectation is, in fact, sufficient for a generalization in probability to hold. Moreover, we will derive a tight concentration bound. Hence, all of the uniform generalization bounds, such as the ones derived in (Alabdulmohsin, 2015), hold, not only in expectation but with a high probability as well. Besides, our result provides, as far as we know, the first strong connection between the two forms of generalization in the literature. We present examples of how our concentration bound can be used to deduce new concentration results for important classes of learning algorithms, such as differential privacy.

The proof of our concentration bound rests on a chain rule that we derive for uniform generalization, which is analogous to the chain rule of mutual information in information theory (Cover and Thomas, 1991). Using the chain rule, we show that learning algorithms that generalize uniformly in expectation are amenable to non-adaptive composition, which is analogous to earlier results using differential privacy, sample compression schemes, and perfect generalization (Dwork and Roth, 2013; Cummings et al., 2016). Moreover, the chain rule lends support to the *information budget* framework, which was recently proposed for controlling the bias of estimators in the adaptive setting using information theory (Russo and Zou, 2016).

The rest of the paper is structured as follows. We will, first, briefly outline the terminology and notation used in this paper and review the existing literature. Next, we recount the main results pertaining to uniform generalization and algorithmic stability and describe how uniform generalization differs from uniform convergence. Finally, we derive the concentration bound for uniform generalization, prove its tightness, and discuss some of its implications afterward.

2 Terminology and Notation

Throughout this paper, we will always write $\mathcal{Z}$ to denote the space of observations (a.k.a. *domain*) and write $\mathcal{H}$ to denote the hypothesis space (a.k.a. *range*). A learning algorithm $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \to \mathcal{H}$ is formally treated as a stochastic map, where the hypothesis $H \in \mathcal{H}$ can be a deterministic or a randomized func-

tion of the sample $S_m \in \mathcal{Z}^m$.

We consider the general setting of learning introduced by Vapnik (Vapnik, 1999). In this setting, the observations $Z_i \in \mathcal{Z}$ can be instance-label pairs $Z_i = (X_i, Y_i)$ as in supervised learning or they can comprise of instances only as in unsupervised learning. The distinction between the two learning paradigms is irrelevant. Moreover, we allow the hypothesis H to be an arbitrary random variable. For instance, H can be a classifier, a regression function, a statistical query, a set of centroids, a density estimate, or an enclosing sphere. Only the relationship between the two random variables S_m and H matters in our analysis.

Moreover, if $Z \sim p(z)$ is a random variable drawn from $\mathcal{Z}$ and $f(Z)$ is a function of Z, we write $\mathbb{E}_{Z\sim p(Z)} f(Z)$ to denote the expected value of $f(Z)$ with respect to the distribution $p(z)$. Occasionally, we omit $p(z)$ from the subscript if it is clear from the context. If Z takes its values from a finite set S uniformly at random, we write $Z \sim S$ to denote this fact. If X is a boolean random variable, then $\mathbb{I}\{X\} = 1$ if and only if X is true, otherwise $\mathbb{I}\{X\} = 0$.

Finally, given two probability measures P and Q defined on the same space, we will write $\langle P,\ Q\rangle$ to denote the *overlapping coefficient* between P and Q. That is, $\langle P,\ Q\rangle = 1 - ||P,\ Q||_{\mathcal{T}}$, where $||P,\ Q||_{\mathcal{T}} = \frac{1}{2}\big|\big|P - Q\big|\big|_1$ is the total variation distance.

3 Related Work

Generalization can be rightfully considered as an extension to the *law of large numbers*, which is one of the earliest and most important results in probability theory and statistics. Suppose we have m i.i.d. observations $S_m = \{Z_1, \ldots, Z_m\} \in \mathcal{Z}^m$ and let $f : \mathcal{Z} \to [0, 1]$ be an arbitrary function. If f is fixed *independently* of S_m, then $\mathbb{E}_{Z_i\sim S_m}[f(Z_i)] \to \mathbb{E}_{Z\sim p(z)}[f(Z)]$ a.s. as $m \to \infty$. This law is generally attributed to Jacob Bernoulli, who wrote an extensive treatise on the subject published posthumously in 1713 (Stigler, 1986). Modern proofs include low-confidence guarantees, e.g. the Chebychev inequality, and high confidence bounds, e.g. the Chernoff method (Boucheron et al., 2004).

When the function f depends on S_m, the law of large numbers is no longer valid because $f(Z_i)$ are not independent random variables. One remedy is to look into the function $F(S_m) = \mathbb{E}_{Z_i\sim S_m} f(Z_i)$. For instance, the Efron-Stein-Steele lemma might be used to bound the variance of F, which, in turn, can be translated into a concentration bound using the Chebychev inequality (Boucheron et al., 2004; Bousquet and Elisseeff, 2002). Alternatively, if F satisfies the *bounded-difference* property, then McDiarmid's inequality yields a high-confidence guarantee (Boucheron et al., 2004; Bousquet and Elisseeff, 2002).

In this paper, the same question is being addressed. However, we address it in an information-theoretic manner. We will show that if the function $f : \mathcal{Z} \to [0, 1]$ (as a random variable instantiated after observing the sample S_m) carries little *information* about any individual observation $Z_i \in S_m$, then the difference between $\mathbb{E}_{Z_i\sim S_m}[f(Z_i)]$ and $\mathbb{E}_{Z\sim p(z)}[f(Z)]$ will be small with a high probability. The measure of information used here is given by the notion of *variational information* $\mathcal{J}(X; Y) = 1 - S(X; Y)$ between the random variables X and Y, where $S(X; Y)$ is the mutual stability introduced in Alabdulmohsin (2015). Variational information is an instance of the class of *informativity* measures using f-divergences, for which an axiomatic basis has been proposed (Csiszár, 1972, 2008).

The information-theoretic approaches of analyzing the generalization risks of learning algorithms, such as the one pursued in this paper, have found applications in adaptive data analysis. This includes the method of Dwork et al. (2015) using the notion of *max-information* and the method of Russo and Zou (2016) using the *mutual information*. For bounded losses, uniform generalization bounds using the *variational information* yield tighter results, as deduced by the Pinsker inequality (Reid and Williamson, 2009). In this paper, we prove that these bounds hold not only in expectation but with a high probability as well.

As a consequence of our main theorem, concentration bounds for a given learning algorithm can be immediately deduced once we recognize that the algorithm generalizes uniformly in expectation. Examples of when this holds include having (1) a finite average description length of the hypothesis, (2) a finite VC dimension of the *induced concept class*, (3) differential privacy, (4) robust generalization, (5) typical generalization, (6) bounded mutual information, and (7) finite domains. We briefly describe these settings that have been previously studied in the literature and prove their connections to uniform generalization in Appendix A. We also present connections between uniform generalization and learnability in Appendix B. A second consequence of our work is establishing the equivalence between the notion of uniform generalization studied by Alabdulmohsin (2015) and the notion of *robust generalization* considered more recently by Cummings et al. (2016).

Besides deriving a concentration bound, we also establish that our bound is tight. This tightness result is inspired by the work of Bassily et al. (2016) (Lemma 7.4) and Shalev-Shwartz et al. (2010) (Example 3), where

similar results are established for differential privacy and learnability respectively. In Section 5.6, we combine techniques from both works to show that our concentration bound is indeed tight.

4 Uniform Generalization

First, we review the main results pertaining to uniform generalization and algorithmic stability. We only mention the key results here for completeness. The reader is referred to Alabdulmohsin (2015) for details.

4.1 Uniform Generalization vs. Uniform Convergence

The main result of Alabdulmohsin (2015) is the *equivalence* between algorithmic stability and *uniform* generalization in expectation across all bounded parametric loss functions that satisfy the Markov chain: $S_m \rightarrow H \rightarrow L(\cdot; H)$.

Definition 3 (Uniform Generalization)**.** *A learning algorithm $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \rightarrow \mathcal{H}$ generalizes uniformly if for any $\epsilon > 0$, $\exists m_0(\epsilon) > 0$ such that for all distributions $p(z)$ on $\mathcal{Z}$, all parametric losses, and all sample sizes $m > m_0(\epsilon)$, we have $|R_{gen}(\mathcal{L})| \leq \epsilon$, where $R_{gen}(\mathcal{L})$ is given in Eq. (2).*

Definition 4. *A learning algorithm $\mathcal{L}$ generalizes uniformly at the rate $\epsilon > 0$ if the expected generalization risk satisfies $|R_{gen}(\mathcal{L})| \leq \epsilon$ for all distributions $p(z)$ on $\mathcal{Z}$ and all parametric losses.*

With some abuse of terminology, we will occasionally say that a learning algorithm generalizes uniformly when it generalizes uniformly according to Definition 4 for some provably small ϵ. Whether we are referring to Definition 3 or 4 will be clear from the context.

Uniform generalization is different from the classical notion of uniform convergence. To see the difference, we note that a parametric loss $L(Z; H) : \mathcal{Z} \times \mathcal{H} \rightarrow [0, 1]$ is a function of the two random variables $Z \in \mathcal{Z}$ and $H \in \mathcal{H}$. This parametric loss on the product space $\mathcal{Z} \times \mathcal{H}$, sometimes called the *loss class* (Bousquet et al., 2004), is a family of loss functions on Z indexed by H. Uniform convergence, such as by using the union bound or the growth function, establishes sufficient conditions for uniform convergence to hold within the family of loss functions indexed by H for *a single* parametric loss. These uniform convergence guarantees are often independent of how $\mathcal{L}$ works.

By contrast, suppose that the learning algorithm $\mathcal{L}$ produces a hypothesis H given a sample $S_m \in \mathcal{Z}^m$ with probability $p_{\mathcal{L}}(H|S_m)$, where $p_{\mathcal{L}}(H|S_m)$ can be degenerate in deterministic algorithms. Then, in principle, one can compute the expected generalization risk $R_{gen}(\mathcal{L})$, defined in Eq. (2), for every possible parametric loss. This is the *average* loss within each possible family of bounded loss functions indexed by H, averaged over the random choice of S_m and the internal randomness of $\mathcal{L}$. Uniform generalization establishes the conditions for $|R_{gen}(\mathcal{L})|$ to go to zero uniformly across all parametric loss functions. It is heavily dependent on how $\mathcal{L}$ works.

4.2 Previous Results

The main result proved in Alabdulmohsin (2015) is that uniform generalization is equivalent to an information-theoretic stability constraint on $\mathcal{L}$.

Definition 5 (Mutual Stability)**.** *The mutual stability between two random variables X and Y is $S(X; Y) \doteq \langle p(X)p(Y),\ p(X, Y)\rangle$, where $\langle P,\ Q\rangle = 1 - ||P\ ,\ Q||_{\mathcal{T}}$, and $||P\ ,\ Q||_{\mathcal{T}}$ is the total variation distance.*

Definition 6 (Variational Information)**.** *The variational information $\mathcal{J}(X; Y)$ between the random variables X and Y is defined by $\mathcal{J}(X; Y) = 1 - S(X; Y)$.*

Informally speaking, $\mathcal{J}(X; Y)$ measures the influence of observing the value of X on the distribution of Y. The rationale behind this definition is revealed next.

Definition 7 (Algorithmic Stability)**.** *Let $\mathcal{L}$ be a learning algorithm that uses $S_m = \{Z_i\}_{i=1,..,m} \sim p^m(z)$ to produce a hypothesis $H \in \mathcal{H}$. Let $Z_{trn} \sim S_m$ be a random variable whose value is drawn uniformly at random from the sample S_m. Then, the algorithmic stability of $\mathcal{L}$ is defined by: $\mathbb{S}(\mathcal{L}) = \inf_{p(z)}\ S(H;\ Z_{trn})$, where the infimum is taken over all possible distributions of observations $p(z)$. A learning algorithm is called stable if $\lim_{m\rightarrow\infty}\ \mathbb{S}(\mathcal{L}) = 1$.*

Intuitively, a learning algorithm is stable if the influence of a *single* training example vanishes as $m \rightarrow \infty$.

Theorem 1 (Alabdulmohsin, 2015)**.** *For any learning algorithm $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \rightarrow \mathcal{H}$, algorithmic stability (Definition 7) is both necessary and sufficient for uniform generalization (Definition 3). In addition, $\left|R_{gen}(\mathcal{L})\right| \leq \mathcal{J}(H;\ Z_{trn}) \leq 1 - \mathbb{S}(\mathcal{L})$, with $R_{gen}(\mathcal{L})$ defined in Eq. (2).*

Theorem 1 reveals that uniform generalization has, at least, three *equivalent* interpretations:

1. *Statistical Interpretation*: A learning algorithm generalizes uniformly if and only if its expected generalization risk $R_{gen}(\mathcal{L})$ vanishes as $m \rightarrow \infty$ uniformly across all bounded parametric losses.

2. *Information-Theoretic Interpretation*: A learning algorithm generalizes uniformly if and only if its

hypothesis H reveals a *vanishing* amount of information about any *single* observation in S_m as $m \to \infty$. This, for example, is satisfied if H has a finite description length or if it is sufficiently randomized as in differential privacy.

3. *Algorithmic Interpreation*: A learning algorithm generalizes uniformly if and only if the contribution of any *single* observation on the hypothesis H vanishes as $m \to \infty$. That is, a learning algorithm generalizes uniformly if and only if it is algorithmically stable.

Other results have also been established in Alabdulmohsin (2015) including the data processing inequality, the information-cannot-hurt inequality, and the uniform generalization bound in the finite hypothesis space setting. Some of those results will be used in our proofs in this paper.

5 Generalization in Expectation vs. Generalization in Probability

The main contribution of this paper is to prove that a uniform generalization in expectation implies a generalization in probability and to derive a tight concentration bound. By contrast, a non-uniform generalization in expectation does not imply that a generalization will actually take place in practice. In addition, we will also establish a *chain rule* for variational information and prove that our large-deviation bound is tight. Interestingly, our proof reveals that uniform generalization is a *robust* property of learning algorithms. Specifically, adding a *finite* amount of information (in bits) to a hypothesis H that generalizes uniformly cannot remove its uniform generalization property.

5.1 Non-Uniform Generalization

We begin by showing why a non-uniform generalization in expectation does not imply concentration[1].

Proposition 1. *There exists a learning algorithm* $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \to \mathcal{H}$ *and a parametric loss* $L(\cdot; H) : \mathcal{Z} \to [0, 1]$ *such that the expected generalization risk is* $R_{gen}(\mathcal{L}) = 0$ *for all* $m \geq 1$*, but for all* $m \geq 1$ *we have* $p\Big\{\big|R_{true}(H) - R_{emp}(H; S_m)\big| = \frac{1}{2}\Big\} = 1$*, where the probability is evaluated over the random choice of* S_m *and the internal randomness of* $\mathcal{L}$*.*

Proposition 1 shows that a generalization in expectation does not imply a generalization in probability. Importantly, it is crucial to observe that the learning algorithm constructed in the proof of Proposition 1 does not, in fact, generalize *uniformly* in expectation. Indeed, this latter observation is not a coincidence as will be proved later in Theorem 4.

5.2 Robustness of Uniform Generalization

Next, we prove that uniform generalization is a robust property of learning algorithms. We will use this fact later to prove that a uniform generalization in expectation implies a generalization in probability. In order to achieve this, we begin with the following chain rule.

Definition 8 (Conditional Variational Information)**.** *The conditional variational information between the two random variables* A *and* B *given* C *is defined by:*

$$\mathcal{J}(A;\, B \,|\, C) = \mathbb{E}_C\big[||p(A, B \,|\, C)\,,\;\; p(A|C) \cdot p(B|C)||_{\mathcal{T}}\big],$$

which is analogous to the notion of conditional mutual information in information theory (Cover and Thomas, 1991).

Theorem 2 (Chain Rule)**.** *Let* $(H_1, \ldots, H_k)$ *be a sequence of random variables. Then, for any random variable* Z*, we have:* $\mathcal{J}(Z; (H_1, ..., H_k)) \leq \sum_{t=1}^{k} \mathcal{J}(Z;\, H_t \,|\, (H_1, ..., H_{t-1}))$

Although the chain rule above provides an upper bound, the upper bound is tight in the following sense:

Proposition 2. *For any random variables* A, B*, and* C*, we have* $\Big|\mathcal{J}(A; (B, C)) - \mathcal{J}(A;\, C \,|\, B)\Big| \leq \mathcal{J}(A;\, B)$ *and* $\Big|\mathcal{J}(A; (B, C)) - \mathcal{J}(A;\, B)\Big| \leq \mathcal{J}(A;\, C \,|\, B)$*.*

In other words, the inequality in the chain rule $\mathcal{J}(A; (B, C)) \leq \mathcal{J}(A;\, B) + \mathcal{J}(A;\, C \,|\, B)$ becomes an equality if $\min\{\mathcal{J}(A;\, B),\; \mathcal{J}(A;\, C \,|\, B)\} = 0$.

The chain rule provides a recipe for computing the bias of estimators when we have a composition of hypotheses $(H_1, \ldots, H_k)$, whether this composition is adaptive or non-adaptive. Recently, Russo and Zou (2016) proposed an *information budget* framework for controlling the bias of estimators by controlling the mutual information between H and the entire sample S_m. The proposed framework rests on the well-known chain rule for mutual information. Here, we note that the chain rule for variational information in Theorem 2 lends further support to the information budget framework.

Next, we use the chain rule in Theorem 2 to prove that uniform generalization is a robust property of learning algorithms. More precisely, if K has a finite domain, then a hypothesis H generalizes uniformly in expectation if and only if the pair (H, K) generalizes uniformly in expectation. Hence, adding any finite amount of information (in bits) to a hypothesis cannot alter its uniform generalization property[2].

[1] Detailed proofs are available in the supplementary file.

[2] Note, by contrast, that the proof of Proposition 1 il-

Theorem 3. *Let $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \rightarrow \mathcal{H}$ be a learning algorithm whose hypothesis is $H \in \mathcal{H}$, which is obtained from a sample S_m. Let $K \in \mathcal{K}$ be a different hypothesis that is obtained from the same sample S_m. If $Z_{trn} \sim S_m$ is a random variable whose value is drawn uniformly at random from S_m, then:*

$$\mathcal{J}(Z_{trn};\, (H, K)) \leq (1 + \frac{|\mathcal{K}|}{2}) \cdot \mathcal{J}(Z_{trn};\, H) + \sqrt{\frac{\log |\mathcal{K}|}{2m}}$$

5.3 Uniform Generalization Implies Concentration

Theorem 3 shows that adding a finite amount of information (in bits) cannot remove the uniform generalization property of learning algorithms. We will use this fact, next, to prove that a uniform generalization in expectation implies a generalization in probability.

The intuition behind the proof is as follows. Suppose we have a hypothesis H that generalizes uniformly in expectation but, for the purpose of obtaining a contradiction, suppose that there exists a parametric loss $L(\cdot; H)$ that does not generalize in probability. Then, *adding* little information to the hypothesis H will allow us to construct a *different* parametric loss that does not generalize in expectation. In particular, we will only to need to know whether the empirical risk w.r.t. $L(\cdot; H)$ is greater than, approximately equal to, or is less than the true risk w.r.t. the same loss. This is described in, at most, two bits. Knowing this additional information, we can define a new parametric loss that does not generalize in expectation, which contradicts the statement of Theorem 3. This line of reasoning is formalized in the following theorem.

Theorem 4. *Let $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \rightarrow \mathcal{H}$ be a learning algorithm, whose risk is evaluated using a parametric loss function $L(\cdot; H) : \mathcal{Z} \rightarrow [0, 1]$. Then:*

$$p\Big\{\big|R_{emp}(H; S_m) - R_{true}(H)\big| \geq t\Big\}$$
$$\leq \frac{5}{2t}\Big[\mathcal{J}(Z_{trn};\, H) + \sqrt{\frac{\log 9}{25m}}\Big] \leq \frac{5}{2t}\Big[1 - \mathbb{S}(\mathcal{L}) + \sqrt{\frac{\log 9}{25m}}\Big],$$

where $\mathbb{S}(\mathcal{L})$ is the algorithmic stability of $\mathcal{L}$ given in Definition 7, and the probability is evaluated over the random choice of S_m and the internal randomness of $\mathcal{L}$. In particular, if $\mathcal{L}$ generalizes uniformly, i.e. $\mathbb{S}(\mathcal{L}) \rightarrow 1$ as $m \rightarrow \infty$, it generalizes in probability for any chosen parametric loss.

The same proof technique used in Theorem 4 also implies the following concentration bound, which is useful when $I(H; S_m) = o(m)$. The following bound compares well with the bound derived in Russo and Zou (2016) using properties of sub-Gaussian loss functions.

Proposition 3. *Let $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \rightarrow \mathcal{H}$ be a learning algorithm, whose risk is evaluated using a parametric loss function $L(\cdot; H) : \mathcal{Z} \rightarrow [0, 1]$. Then:*

$$p\Big\{\big|R_{emp}(H; S_m) - R_{true}(H)\big| \geq t\Big\} \leq \frac{1}{t}\sqrt{\frac{I(S_m; H) + 3}{2m}},$$

Note that having a bounded mutual information, i.e. $I(S_m; H) = o(m)$, which is the setting recently considered in the work of Russo and Zou (2016), is sufficient for uniform generalization to hold.

5.4 Implications

5.4.1 Concentration

In Alabdulmohsin (2015), it was shown that the notion of uniform generalization allows us to reason about learning algorithms in pure information-theoretic terms. This is because uniform generalization is equivalent to an information-theoretic algorithmic stability constraint on learning algorithms. For example, the data processing inequality implies that one can improve the uniform generalization risk by either post-processing the hypothesis, such as sparsification or decision tree pruning, or by pre-processing training examples, such as by introducing noise, Tikhonov regularization, or dropout. Needless to mention, both are common techniques in machine learning. In addition, uniform generalization recovers classical results in learning theory, such as the generalization bounds in the finite hypothesis space setting and finite domains (Alabdulmohsin, 2015). However, such conclusions previously held only *in expectation*.

The most important implication of Theorem 4 is to establish that such conclusions actually hold with a high probability as well. In addition, the concentration bound derived in Theorem 4 shows that algorithmic stability $\mathbb{S}(\mathcal{L})$ not only controls the generalization risk of $\mathcal{L}$ in expectation, i.e. due to its equivalence with uniform generalization, but it also controls the rate of convergence in probability. This brings us to the following important remark:

Remark 1. *By improving algorithmic stability, we improve both the expectation of the generalization risk* and *its variance.*

Besides, Theorem 4 can be useful in deriving new concentrations bounds for important classes of learning algorithms once we recognize the existence of uniform generalization. We illustrate this technique on *differential privacy* next.

lustrates an example where a hypothesis H may generalize perfectly in expectation w.r.t. a fixed parametric loss, but a single bit of information suffices to destroy this generalization advantage. This never occurs when H generalizes uniformly since uniform generalization is a robust property.

5.5 Differential Privacy

Differential privacy addresses the goal of obtaining useful information about the sample S_m as a whole without revealing a lot of information about each individual observation in the sample (Dwork and Roth, 2013). It closely resembles the notion of algorithmic stability proposed in Alabdulmohsin (2015) because a learning algorithm is stable according to the latter definition if and only if the posterior distribution of an individual observation Z_{trn} in the sample S_m becomes arbitrarily close, in the total variation distance, to the prior distribution $p(Z_{trn})$ as $m \to \infty$. Indeed, differential privacy is a stronger privacy guarantee.

Definition 9 (Dwork & Roth, 2013)**.** *A randomized learning algorithm $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \to \mathcal{H}$ is (ϵ, δ) differentially private if for any $\mathcal{O} \subseteq \mathcal{H}$ and any two samples S and S' that differ in one observation only, we have:*

$$p(H \in \mathcal{O} \mid S) \leq e^{\epsilon} \cdot p(H \in \mathcal{O} \mid S') + \delta$$

Concentration bounds for differential privacy have been derived, such as the recent work of Bassily et al. (2016). Nevertheless, we remark here that Theorem 4 can be used to derive a new concentration bound for differential privacy. Comparing our bound with the lower bound of Lemma 7.4 in Bassily et al. (2016) reveals that the dependence on δ and t is tight up to a constant factor.

Corollary 1. *If a learning algorithm $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \to \mathcal{H}$ is (ϵ, δ) differentially private, then: $p\Big\{\big|R_{emp}(H; S_m) - R_{true}(H)\big| \geq t\Big\} \leq \frac{5}{4t}\Big[e^{\epsilon} - 1 + \delta + \sqrt{\frac{2 \log 9}{25m}}\Big]$.*

Not surprisingly, the differential privacy parameters (ϵ, δ) control the generalization risk of differential privacy, with the quantity $(e^{\epsilon} - 1 + \delta)$ acting a role that is analogous to the standard error.

5.5.1 Equivalnce with Robust Generalization

Another implication of the concentration bound in Theorem 4 is establishing the *equivalence* between the notion of uniform generalization and the notion of *robust generalization* studied in Cummings et al. (2016).

Definition 10 (Robust Generalization)**.** *A learning algorithm $\mathcal{L}$ is (ϵ, δ) robustly generalizing if for all distribution $p(z)$ on $\mathcal{Z}$ and any binary-valued parametric loss function $L(\cdot; H) : \mathcal{Z} \to \{0, 1\}$ that satisfies the Markov chain $S_m \to H \to L(\cdot; H)$, we have with a probability of at least $1 - \zeta$ over the choice of S that:*

$$p\Big\{\big|\mathbb{E}_{Z \sim p(z)} L(Z; H) - \frac{1}{m} \sum_{Z_i \in S} L(Z_i; H)\big| \leq \epsilon\Big\} \geq 1 - \gamma,$$

for some γ, ζ such that $\delta = \gamma + \zeta$[3]*.*

In the following theorem, we prove that robust generalization is equivalent to uniform generalization.

Corollary 2. *If a learning algorithm $\mathcal{L}$ is (ϵ, δ) robustly generalizing, then it generalizes uniformly at the rate $\epsilon + \delta$. Conversely, if a learning algorithm generalizes uniformly with rate τ, then it is (ϵ, γ) robustly generalizing with $\gamma = (5/2)(\tau + \sqrt{\log 9/(25m)})/\epsilon$. Moreover, if $\mathbb{S}(\mathcal{L}) \to 0$ as $m \to \infty$, then both γ and ϵ can be made arbitrarily close to zero using a sufficiently large sample size m.*

5.6 Tightness Result

Finally,we note that the concentration bound has a linear dependence on the algorithmic stability term $1 - \mathbb{S}(\mathcal{L})$ or, in a distribution-dependent manner, on the variational information $\mathcal{J}(Z_{trn}; H)$. Typically, $\mathcal{J}(Z_{trn}; H) = \Theta(1/\sqrt{m})$. By contrast, the VC bound provides an exponential decay for supervised classification tasks (Vapnik, 1999; Shalev-Shwartz et al., 2010). This raises the question of whether or not the concentration bound in Theorem 4 can be improved. In this section, we prove that the bound is actually tight. The following theorem is inspired by the work of (Bassily et al., 2016) (Lemma 7.4) and Shalev-Shwartz et al. (2010) (Example 3), who established similar results for differential privacy and learnability respectively.

Theorem 5. *For any rational $0 < t < 1$, there exists a learning algorithm $\mathcal{L} : \cup_{m=1}^{\infty} \mathcal{Z}^m \to \mathcal{H}$, a sample size m, a distribution $p(z)$, and a parametric loss $L(\cdot; H) : \mathcal{Z} \to [0, 1]$ such that $\mathcal{L}$ generalizes uniformly in expectation and it satisfies:*

$$p\Big\{\big|R_{emp}(H; S_m) - R_{true}(H)\big| = t\Big\} = \frac{\mathcal{J}(Z_{trn}; H)}{t}$$

Theorem 5 shows that, without making any additional assumptions beyond that of uniform generalization, the concentration bound in Theorem 4 is tight up to constant factors. Essentially, the only difference between the upper and the lower bounds is a vanishing $O(1/\sqrt{m})$ term that is *independent* of $\mathcal{L}$.

6 Conclusions

Uniform generalization in expectation is a notion of generalization that is equivalent to an information-theoretic algorithmic stability constraint on learning

[3]The original definition proposed in Cummings et al. (2016) states that the probability is evaluted over any "adversary" that takes the hypothesis H as input to produce a loss function $L(\cdot; H)$. However, this is equivalent to the Markov chain $S_m \to H \to L(\cdot; H)$.

algorithms. In this paper, we proved that whereas generalization in expectation does not imply a generalization in probability, a uniform generalization in expectation implies a generalization in probability and we derived a tight concentration bound. The bound reveals that algorithmic stability improves both the expectation of the generalization risk and its variance. Hence, by constraining the "amount" of information that a hypothesis can carry about any *individual* training example or, equivalently, by limiting the "size" of the contribution of any individual training example on the final hypothesis, the learning algorithm is guaranteed to generalize well with a high probability. Furthermore, we proved a chain rule for variational information, which revealed that uniform generalization is a robust property of learning algorithms. Finally, we proved that the concentration bound is tight.

A Relations to Other Notions of Generalization & Stability

The connection between differential privacy and uniform generalization is summarized as follows.

Proposition 4. *Let $\mathcal{L}$ be a (ϵ, δ)-differentially private learning algorithm. Let $Z_{trn} \sim S$ be a single training example and let $H \sim p_{\mathcal{L}}(h|S)$ be the hypothesis produced by $\mathcal{L}$. Then $\mathcal{J}(Z_{trn};\, H) \leq \frac{e^{\epsilon}-1+\delta}{2}$.*

Next, it can be shown that perfect generalization implies differential privacy (Cummings et al., 2016) so it implies uniform generalization. Also, sample compression implies robust generalization (Cummings et al., 2016), which in turn implies a uniform generalization. Moreover, typical stability (Bassily and Freund, 2016) is equivalent to perfect generalization when the observations are drawn i.i.d., so it implies uniform generalization as well.

The proof that a bounded mutual information, i.e. having $I(S_m; H) = o(m)$, implies uniform generalization is a direct consequence of the concentration bound in Proposition 3.

Finally, the proofs that a finite hypothesis space, a finite VC dimension in the induced concept class, and a finite domain are each sufficient for uniform generalization to hold are provided in Alabdulmohsin (2015).

B Uniform Generalization and Learnability

B.1 Consistency of Empirical Risk Minimization

Uniform generalization is a sufficient condition for the consistency of empirical risk minimization (ERM). Suppose we have an ERM learning algorithm, whose hypothesis is denoted H_{ERM}. We have by definition:

$$\begin{aligned} R_{emp}(\mathcal{L}) &= \mathbb{E}_S \mathbb{E}_{H|S}[R_{emp}(H;S)] = \mathbb{E}_S[\min_{h\in\mathcal{H}} R_{emp}(h;S)] \\ &\leq \min_{h\in\mathcal{H}} \big[\mathbb{E}_S R_{emp}(h;S)\big] = \min_{h\in\mathcal{H}} R_{true}(h) \\ &= R_{true}(h^{\star}), \end{aligned}$$

where $h^{\star}$ is the optimal hypothesis. However, the true risk of $\mathcal{L}$ satisfies:

$$R_{true}(\mathcal{L}) - R_{true}(h^{\star}) \leq R_{true}(\mathcal{L}) - R_{emp}(\mathcal{L})$$

Thus, algorithmic stability of ERM implies consistency because $\mathcal{J}(Z_{trn};\, H_{ERM})$ provides an upper bound on $|R_{true}(\mathcal{L}) - R_{emp}(\mathcal{L})|$. In fact, because $R_{true}(H_{ERM}) - R_{true}(h^{\star}) \geq 0$, we have by the Markov inequality:

$$p\Big\{R_{true}(H_{ERM}) - R_{true}(h^{\star}) \geq t\Big\} \leq \frac{\mathcal{J}(Z_{trn};\, H_{ERM})}{t}$$

B.2 Sample Compression and Learnability

Moreover, recent results on the connection between sample compression schemes and learnability (David et al., 2016) reveal that any learnable hypothesis space is learnable using an algorithm that generalizes uniformly in expectation, with only a logarithmic increase in the sample complexity. Because sample compression schemes satisfy robust generalization (Cummings et al., 2016), they generalize uniformly in expectation.

C Implications of the Chain Rule

The chain rule in Theorem 2 along with the fact that for any random variables X, Y, Z, we have $\mathcal{J}(X;\, Y) \leq \mathcal{J}(X;\, Y, Z)$ can both be used to derive many results. We illustrate this with two examples here. First, suppose we have the Markov chain $A \to B \to C$. By the chain rule, we have:

$$\mathcal{J}(A;\, C) \leq \mathcal{J}(A;\, (B, C)) \leq \mathcal{J}(A;\, B) + \mathcal{J}(A;\, C \,|\, B)$$

Since $A \to B \to C$ implies $\mathcal{J}(A;\, C \,|\, B) = 0$ by the Markov property, we conclude that $\mathcal{J}(A;\, C) \leq \mathcal{J}(A;\, B)$. This is the *data processing* inequality. In general, we have the triangle-like inequality:

$$\mathcal{J}(X;\, Y) \leq \mathcal{J}(X;\, Z) + \mathcal{J}(X;\, Y \,|\, Z) \tag{3}$$

Second, if $A \to B \to C$, we have:

$$\mathcal{J}(A;\, B) \leq \mathcal{J}(A;\, (B, C)) \leq \mathcal{J}(A;\, B) + \mathcal{J}(A;\, C \,|\, B)$$

Since $\mathcal{J}(A;\, C \,|\, B) = 0$, we conclude that $A \to B \to C$ implies $\mathcal{J}(A;\, (B, C)) = \mathcal{J}(A;\, B)$. Both results were used previously, and were proved using different methods in Alabdulmohsin (2015).

D Proof of Proposition 1

Let $\mathcal{Z} = [0,1]$ be an instance space with a continuous marginal density $p(z)$ (hence, has no atoms) and let $\mathcal{Y} = \{-1,+1\}$ be the target set. Let $h^\star : \mathcal{Z} \to \{-1,+1\}$ be some *fixed* predictor, such that $p\{h^\star(Z)=1\} = \frac{1}{2}$, where the probability is evaluated over the random choice of $Z \in \mathcal{Z}$. In other words, the marginal distribution of the labels predicted by $h^\star(\cdot)$ is uniform[4].

Next, let the hypothesis space $\mathcal{H}$ be the set of predictors from $\mathcal{Z}$ to $\{-1,+1\}$ that output a label in $\{-1,+1\}$ uniformly at random everywhere in $\mathcal{Z}$ except at a finite number of points. Therefore, the hypothesis $H : \mathcal{Z} \to \{-1,+1\}$ selected by the learning algorithm is a predictor. Define the parametric loss by $L(Z;H) = \mathbb{I}\big\{H(Z) \neq h^\star(Z)\big\}$.

Next, we construct a learning algorithm $\mathcal{L}$ that generalizes perfectly in expectation but it does not generalize in probability. The learning algorithm $\mathcal{L}$ simply picks one of $H^0_{S_m}(\cdot)$ or $H^1_{S_m}(\cdot)$ with equal probability, where:

$$H^0_{S_m}(Z) = \begin{cases} -h^\star(Z) & \text{if } Z \in S_m \\ \text{Uniform}(-1,+1) & \text{if } Z \notin S_m \end{cases}$$

$$H^1_{S_m}(Z) = \begin{cases} h^\star(Z) & \text{if } Z \in S_m \\ \text{Uniform}(-1,+1) & \text{if } Z \notin S_m \end{cases}$$

Because $\mathcal{Z}$ is uncountable, where the probability of seeing the same observation Z twice is zero, $R_{true}(H) = \frac{1}{2}$ for this learning algorithm. Thus:

$$R_{gen}(\mathcal{L}) = \mathbb{E}_{S_m,H}\big[R_{emp}(H;S_m) - R_{true}(H)\big] = 0$$

However, the empirical risk for any S_m satisfies $R_{emp}(H;S_m) \in \{0,1\}$ while the true risk always satisfies $R_{true}(H) = \frac{1}{2}$, as mentioned earlier. Hence, the statement of the proposition follows.

Finally, we prove that the algorithm does not generalize uniformly in expectation. There are, at least, two ways of showing this. The first approach is to use the equivalence between uniform generalization and algorithmic stability as stated in Theorem 1. Given the hypothesis $H \in \{H^0_{S_m},\ H^1_{S_m}\}$ learned by the algorithm constructed here, the marginal distribution of an individual training example $p(Z_{trn}|H)$ is uniform over the sample S_m. This follows from the fact that the hypothesis H has to encode the entire sample S_m. However, the probability of seeing the same observation twice is zero (by construction). Hence, $||p(Z_{trn})\,,\ \ p(Z_{trn}|H)||_\mathcal{T} = 1$ for all H. This shows that $\mathbb{S}(\mathcal{L}) = 0$ for all $m \geq 1$, and the learning algorithm is not stable. Therefore, by Theorem 1, it does not generalize uniformly. Note that we used the information-theoretic interpretation of uniform generalization.

The second approach is to use the statistical interpretation of uniform generalization. Let $H \in \{H^0_{S_m},\ H^1_{S_m}\}$ be the hypothesis inferred by the learning algorithm above, and consider the following *different* parametric loss:

$$L(Z;H^k_{S_m}) = \mathbb{I}\big\{(-1)^{k+1}H^k_{S_m}(Z) \neq h^\star(Z)\big\}$$

In other words, we flip the predictions of $H^k_{S_m}$ if $k=0$ and measure the misclassification loss afterwards. Note that this is a parametric loss; it has a bounded range and satisfies the Markov chain $S_m \to H \to L(\cdot;H)$. However, the expected generalization risk w.r.t. this parametric loss is $R_{gen}(\mathcal{L}) = \frac{1}{2}$ for all $m \geq 1$ because $R_{emp}(\mathcal{L}) = 0$ w.r.t. to this loss. Therefore, $\mathcal{L}$ does not generalize uniformly in expectation[5].

E Proof of Theorem 2

We will first prove the inequality when $k=2$. First, we write by definition:

$$\mathcal{J}(Z;\,(H_1,H_2)) = ||p(Z,H_1,H_2)\,,\ \ p(Z)\,p(H_1,H_2)||_\mathcal{T}$$

Using the fact that the total variation distance is related to the ℓ_1 distance by $||P\,,\ \ Q||_\mathcal{T} = \frac{1}{2}||P-Q||_1$, we have:

$$\begin{aligned}
\mathcal{J}(Z;\,(H_1,H_2)) &= \frac{1}{2}\Big|\Big|\ p(Z,H_1,H_2) - p(Z)\,p(H_1,H_2)\ \Big|\Big|_1 \\
&= \frac{1}{2}\Big|\Big|\ p(Z,H_1)\,p(H_2|Z,H_1) - p(Z)\,p(H_1)\,p(H_2|H_1)\ \Big|\Big|_1 \\
&= \frac{1}{2}\Big|\Big|\ \big[p(Z,H_1) - p(Z)\,p(H_1)\big]\cdot p(H_2|H_1) \\
&\qquad + p(Z,H_1)\ \cdot \big[p(H_2|Z,H_1) - p(H_2|H_1)\big]\ \Big|\Big|_1
\end{aligned}$$

Using the triangle inequality:

$$\begin{aligned}
&\mathcal{J}(Z_{trn};\,(H_1,H_2)) \\
&\quad \leq \frac{1}{2}\Big|\Big|\big[p(Z,H_1) - p(Z)\,p(H_1)\big]\cdot p(H_2|H_1)\Big|\Big|_1 \\
&\qquad + \frac{1}{2}\Big|\Big|p(Z,H_1)\cdot\big[p(H_2|Z,H_1) - p(H_2|H_1)\big]\Big|\Big|_1
\end{aligned}$$

The above inequality is interpreted by expanding the ℓ_1 distance into a sum of absolute values of terms in

[4] These assumptions are satisfied, for example, if $p(z)$ is uniform in $[0,1]$ and $h^\star(Z) = \mathbb{I}\{Z < \frac{1}{2}\}$.

[5] We remark here that the Markov inequality cannot be used to provide a concentration bound for the learning algorithm in Proposition 1 even if the expected generalization risk goes to zero because the quantity $R_{emp}(H;S_m) - R_{true}(H)$ is not guaranteed to be non-negative. Indeed, this is precisely why the learning algorithm $\mathcal{L}$ constructed here generalizes in expectation but not in probability.

the product space $\mathcal{Z} \times \mathcal{H}_1 \times \mathcal{H}_2$, where $H_k \in \mathcal{H}_k$. Next, we bound each term on the right-hand side separately. For the first term, we note that:

$$\begin{aligned}\frac{1}{2}\big|\big|\ \big[p(Z, H_1) - p(Z)\, p(H_1)\big] \cdot p(H_2|H_1)\ \big|\big|_1 \\ = \frac{1}{2}\big|\big|\ p(Z, H_1) - p(Z)\, p(H_1)\ \big|\big|_1 = \mathcal{J}(Z;\, H_1)\end{aligned} \tag{4}$$

The equality holds by expanding the ℓ_1 distance and using the fact that $\sum_{H_2} p(H_2|H_1) = 1$.

However, the second term can be re-written as:

$$\begin{aligned}&\frac{1}{2}\big|\big|\ p(Z, H_1)\ \cdot \big[p(H_2|Z, H_1) - p(H_2|H_1)\big]\ \big|\big|_1 \\ &= \frac{1}{2}\big|\big|\ p(H_1) \cdot \big[p(H_2, Z|H_1) - p(Z|H_1)\, p(H_2|H_1)\big]\ \big|\big|_1 \\ &= \mathbb{E}_{H_1}\big[||p(H_2, Z|H_1)\,,\ \ p(Z|H_1)\, p(H_2|H_1)||_{\mathcal{T}}\big] \\ &= \mathcal{J}(Z;\, H_2 \,|\, H_1)\end{aligned} \tag{5}$$

Combining Eq. (4) and (5) yields the inequality:

$$\mathcal{J}(Z;\, (H_1, H_2)) \leq \mathcal{J}(Z;\, H_1) + \mathcal{J}(Z;\, H_2 \,|\, H_1) \tag{6}$$

Next, we use Eq. (6) to prove the general statement for all $k \geq 1$. By writing:

$$\begin{aligned}\mathcal{J}(Z;\, (H_1, \ldots, H_k)) \leq \mathcal{J}(Z;\, H_k \,|\, (H_1, \ldots, H_{k-1})) \\ + \mathcal{J}(Z;\, (H_1, \ldots, H_{k-1}))\end{aligned}$$

Repeating the same inequality on the last term on the right-hand side yields the statement of the theorem.

F Proof of Proposition 2

We will use the following fact (Alabdulmohsin, 2015):

Fact 1 (Information Cannot Hurt)**.** *For any random variables X, Y, Z:*

$$\mathcal{J}(X;\, Y) \leq \mathcal{J}(X;\, (Y, Z))$$

Now, by the triangle inequality:

$$\begin{aligned}\mathcal{J}(A;\, C \,|\, B) &= \mathbb{E}_B||p(A|B) \cdot p(C|B)\,,\ \ p(A, C|B)||_{\mathcal{T}} \\ &= \mathbb{E}_{A,B}||p(C|B)\,,\ \ p(C|A, B)||_{\mathcal{T}} \\ &\leq \mathbb{E}_{A,B}||p(C|B)\,,\ \ p(C)||_{\mathcal{T}} \\ &\quad + \mathbb{E}_{A,B}||p(C)\,,\ \ p(C|A, B)||_{\mathcal{T}} \\ &= \mathbb{E}_B||p(C|B)\,,\ \ p(C)||_{\mathcal{T}} \\ &\quad + \mathbb{E}_{A,B}||p(C)\,,\ \ p(C|A, B)||_{\mathcal{T}} \\ &= \mathcal{J}(B;\, C) + \mathcal{J}(C;\, (A, B))\end{aligned}$$

Therefore:

$$\mathcal{J}(C;\, (A, B)) \geq \mathcal{J}(A;\, C \,|\, B) - \mathcal{J}(B;\, C)$$

Combining this with the following chain rule of Theorem 2:

$$\mathcal{J}(C;\, (A, B)) \leq \mathcal{J}(A;\, C \,|\, B) + \mathcal{J}(B;\, C)$$

yields:

$$\Big|\mathcal{J}(C;\, (A, B)) - \mathcal{J}(A;\, C \,|\, B)\Big| \leq \mathcal{J}(B;\, C)$$

Or equivalently:

$$\Big|\mathcal{J}(A;\, (B, C)) - \mathcal{J}(A;\, C \,|\, B)\Big| \leq \mathcal{J}(A;\, B) \tag{7}$$

To prove the other inequality, we use Fact 1. We have:

$$\mathcal{J}(A;\, B) \leq \mathcal{J}(A;\, (B, C)) \leq \mathcal{J}(A;\, B) + \mathcal{J}(A;\, C \,|\, B),$$

where the first inequality follows from Fact 1 and the second inequality follows from the chain rule. Thus, we obtain the desired bound:

$$\Big|\mathcal{J}(A;\, (B, C)) - \mathcal{J}(A;\, B)\Big| \leq \mathcal{J}(A;\, C \,|\, B) \tag{8}$$

Both Eq. 7 and Eq. 8 imply that the chain rule is tight. More precisely, the inequality can be made arbitrarily close to an equality when one of the two terms in the upper bound is chosen to be arbitrarily close to zero.

G Proof of Theorem 3

We will use the following fact:

Fact 2. *Let $f : \mathcal{X} \to [0, 1]$ be a function with a bounded range in the interval $[0, 1]$. Let $p_1(x)$ and $p_2(x)$ be two different probability measures defined on the same space $\mathcal{X}$. Then:*

$$\Big|\mathbb{E}_{X \sim p_1(x)} f(X) - \mathbb{E}_{X \sim p_2(x)} f(X)\Big| \leq ||p_1(x)\,,\ \ p_2(x)||_{\mathcal{T}}$$

First, consider the following scenario. Suppose a learning algorithm $\mathcal{L}$ generates a hypothesis $H \in \mathcal{H}$ from some marginal distribution $p(h)$ *independently* of the sample S_m. Afterward, a sample $S_m \in \mathcal{Z}^m$ is observed, which comprises of m i.i.d. observations. Then, $\mathcal{L}$ selects $K \in \mathcal{K}$ according to $p(k|H, S_m)$.

In this scenario, we have:

$$\mathcal{J}(Z_{trn};\, (H, K)) = \mathcal{J}(Z_{trn};\, K \,|\, H),$$

where the equality follows from the chain rule in Theorem 2, the statement of Proposition 2, and the fact that $\mathcal{J}(Z_{trn};\, H) = 0$. The conditional variational information is written as:

$$\begin{aligned}&\mathcal{J}(Z_{trn};\, K \,|\, H) \\ &\quad = \mathbb{E}_H||p(Z_{trn}) \cdot p(K|H)\,,\ \ p(Z_{trn}, K|H)||_{\mathcal{T}},\end{aligned}$$

where we used the fact that $p(Z_{trn}|H) = p(Z_{trn})$. Next, by marginalization, the conditional distribution $p(K|H)$ is given by:

$$\begin{aligned} p(K|H) &= \mathbb{E}_{Z'_{trn}|H}[p(K|Z'_{trn}, H)] \\ &= \mathbb{E}_{Z'_{trn}}[p(K|Z'_{trn}, H)]. \end{aligned}$$

where the expectation is taken with respect to the marginal distribution of observations $p(z)$. Similarly:

$$\begin{aligned} p(Z_{trn}, K|H) &= p(Z_{trn}|H) \cdot p(K|Z_{trn}, H) \\ &= p(Z_{trn}) \cdot p(K|Z_{trn}, H) \end{aligned}$$

Therefore:

$$\begin{aligned} &\mathcal{J}(Z_{trn};\, K \,|\, H) \\ &\quad = \mathbb{E}_H \mathbb{E}_{Z_{trn}} || \mathbb{E}_{Z'_{trn}} p(K|Z'_{trn}, H)\,,\ p(K|Z_{trn}, H) ||_{\mathcal{T}} \end{aligned}$$

Next, for every value of H that is generated independently of the sample S_m, the variational information between $Z_{trn} \sim S_m$ and $K \in \mathcal{K}$ can be bounded using Theorem 3 in (Alabdulmohsin, 2015). This follows because H is selected independently of the sample S_m, and, hence, the i.i.d. property of the observations Z_i continue to hold. Therefore, we obtain:

$$\begin{aligned} &\mathbb{E}_H \mathbb{E}_{Z_{trn}} || \mathbb{E}_{Z'_{trn}} p(K|Z'_{trn}, H)\,,\ p(K|Z_{trn}, H) ||_{\mathcal{T}} \\ &\qquad = \mathcal{J}(Z_{trn};\, K \,|\, H) \\ &\qquad \leq \sqrt{\frac{|\mathcal{K}|}{2m}} \end{aligned} \tag{9}$$

Because $p(K|Z_{trn}, H)$ is arbitrary, the above bound holds for any distribution of observations $p(z)$, any distribution $p(h)$, and any family of conditional distributions $p(k|Z_{trn}, H)$.

Next, we return to the original setting where both $H \in \mathcal{H}$ and $K \in \mathcal{K}$ are chosen according to the sample S_m. We have:

$$\begin{aligned} &\mathcal{J}(Z_{trn};\, K \,|\, H) \\ &\quad = \mathbb{E}_H || p(Z_{trn}|H) \cdot p(K|H)\,,\ p(Z_{trn}, K|H) ||_{\mathcal{T}} \\ &\quad = \mathbb{E}_{H, Z_{trn}} || p(K|H)\,,\ p(K|Z_{trn}, H) ||_{\mathcal{T}} \\ &\quad = \mathbb{E}_{H, Z_{trn}} || \mathbb{E}_{Z'_{trn}|H}[p(K|Z'_{trn}, H)]\,,\ p(K|Z_{trn}, H) ||_{\mathcal{T}} \\ &\quad \leq \mathbb{E}_{H, Z_{trn}} || \mathbb{E}_{Z'_{trn}|H}[p(K|Z'_{trn}, H)]\,,\ \mathbb{E}_{Z'_{trn}}[p(K|Z'_{trn}, H)] ||_{\mathcal{T}} \\ &\qquad + \mathbb{E}_{H, Z_{trn}} \Big|\Big| \mathbb{E}_{Z'_{trn}}[p(K|Z'_{trn}, H)] - p(K|Z_{trn}, H) \Big|\Big|_1 \end{aligned} \tag{10}$$

We bound the second term in Eq. 10 using our earlier result in Eq. 9.

Next, we would like to bound the first term. Using the fact that the total variation distance is related to the ℓ_1 distance by $||P\,,\ Q||_{\mathcal{T}} = \frac{1}{2}||P - Q||_1$, we have:

$$\begin{aligned} &\mathbb{E}_{H, Z_{trn}} || \mathbb{E}_{Z'_{trn}|H}[p(K|Z'_{trn}, H)]\,,\ \mathbb{E}_{Z'_{trn}}[p(K|Z'_{trn}, H)] ||_{\mathcal{T}} \\ &= \mathbb{E}_H \left|| \mathbb{E}_{Z'_{trn}|H}[p(K|Z'_{trn}, H)]\,,\ \mathbb{E}_{Z'_{trn}}[p(K|Z'_{trn}, H)] ||_{\mathcal{T}} \right. \\ &= \frac{1}{2}\, \mathbb{E}_H \sum_{K \in \mathcal{K}} \Big| \mathbb{E}_{Z'_{trn}|H}[p(K|Z'_{trn}, H)] - \mathbb{E}_{Z'_{trn}}[p(K|Z'_{trn}, H)] \Big| \\ &\leq \frac{1}{2}\, \mathbb{E}_H \sum_{K \in \mathcal{K}} || p(Z'_{trn}|H)\,,\ p(Z'_{trn}) ||_{\mathcal{T}} \\ &= \frac{1}{2} \sum_{K \in \mathcal{K}} \mathbb{E}_H || p(Z'_{trn}|H)\,,\ p(Z'_{trn}) ||_{\mathcal{T}} \\ &= \frac{|\mathcal{K}|}{2}\, \mathcal{J}(Z_{trn};\, H) \end{aligned} \tag{11}$$

Here, the inequality follows from Fact 2.

Combining all results in Eq. 9, 10, and 11:

$$\mathcal{J}(Z_{trn};\, K \,|\, H) \leq \frac{|\mathcal{K}|}{2}\, \mathcal{J}(Z_{trn};\, H) + \sqrt{\frac{|\mathcal{K}|}{2m}} \tag{12}$$

This along with the chain rule imply the statement of the theorem.

H Proof of Theorem 4

Let $L(\cdot; H)$ be a parametric loss function and write:

$$\kappa(t) = p\Big\{ \big| R_{emp}(H; S_m) - R_{true}(H) \big| \geq t \Big\} \tag{13}$$

Consider the new pair of hypotheses (H, K), where:

$$K = \begin{cases} +1, & \text{if } R_{emp}(H; S_m) \geq R_{true}(H) + t \\ -1, & \text{if } R_{emp}(H; S_m) \leq R_{true}(H) - t \\ 0, & \text{otherwise} \end{cases}$$

Then, by Theorem 3, the uniform generalization risk in expectation for the composition of hypotheses (H, K) is bounded by $4(5/2)\, \mathcal{J}(Z_{trn};\, H) + \sqrt{\frac{\log 3}{2m}}$. This holds uniformly across all parametric loss functions $L'(\cdot; H, K) \rightarrow [0, 1]$ that satisfy the Markov chain $S_m \rightarrow (H, K) \rightarrow L'(\cdot; H, K)$. Next, consider the parametric loss:

$$L'(Z; H, K) = \begin{cases} L(Z; H) & \text{if } K = +1 \\ 1 - L(Z; H) & \text{if } K = -1 \\ 0 & \text{otherwise} \end{cases}$$

Note that $L'(\cdot; H, K)$ is parametric with respect to the composition of hypotheses (H, K). Using Eq. 13, the generalization risk w.r.t $L'(\cdot; H, K)$ in expectation is, at least, as large as $t\, \kappa(t)$. Therefore, by Theorem 1 and Theorem 3, we have $t\, \kappa(t) \leq (5/2)\, \mathcal{J}(Z_{trn};\, H) + \sqrt{\frac{\log 3}{2m}}$. Because $\mathcal{J}(Z_{trn};\, H) \leq 1 - \mathbb{S}(\mathcal{L})$ by definition, the statement of the theorem immediately follows.

I Proof of Proposition 3

Let $I(X;Y)$ denote the mutual information between X and Y and let $\mathbf{H}(X)$ denote the Shannon entropy of the random variable X measured in nats (i.e. using natural logarithms). We write $S_m = (Z_1, \ldots, Z_m)$. We have:

$$\begin{aligned} I(S_m;(H,K)) &= \mathbf{H}(S_m) - \mathbf{H}(S_m \mid H,K) \\ &= \sum_{i=1}^{m} \mathbf{H}(Z_i) - \sum_{i=1}^{m} \mathbf{H}(Z_i|H,K,Z_1,\ldots,Z_{i-1}) \\ &\geq \sum_{i=1}^{m} \mathbf{H}(Z_i) - \mathbf{H}(Z_i|H,K) \\ &= mI(Z_{trn};H,K) \end{aligned}$$

The second line is the chain rule for entropy and the third lines follows from the fact that conditioning reduces entropy. We obtain:

$$I(Z_{trn};H,K) \leq \frac{I(S_m;(H,K))}{m}$$

By Pinsker's inequality:

$$\mathcal{J}(Z_{trn};\,(H,K)) \leq \sqrt{\frac{I(Z_{trn};(H,K))}{2}} \leq \sqrt{\frac{I(S_m;(H,K))}{2m}}$$

Using the chain rule for mutual information:

$$\begin{aligned} \mathcal{J}(Z_{trn};\,(H,K)) &\leq \sqrt{\frac{I(S_m;(H,K))}{2m}} \\ &= \sqrt{\frac{I(S_m;H) + I(S_m;K|H)}{2m}} \\ &\leq \sqrt{\frac{I(S_m;H) + \mathbf{H}(K)}{2m}} \\ &\leq \sqrt{\frac{I(S_m;H) + \log|\mathcal{K}|}{2m}} \end{aligned}$$

The desired bound follows by applying the same proof technique of Theorem 4 on the last uniform generalization bound.

J Proof of Corollary 1

First, we note that for any two adjacent samples S and S', we have:

$$p(H|S) - p(H|S') \leq (e^{\epsilon} - 1)\, p(H|S') + \delta$$

This follows by definition of differential privacy. Similarly, we have:

$$\begin{aligned} p(H|S) - p(H|S') &\geq (e^{-\epsilon} - 1)\, p(H|S') - e^{-\epsilon}\delta \\ &= -\Big[(1 - e^{-\epsilon})\, p(H|S') + e^{-\epsilon}\delta\Big] \\ &\geq -e^{\epsilon}\,\Big[(1 - e^{-\epsilon})\, p(H|S') + e^{-\epsilon}\delta\Big] \\ &= -\Big[(e^{\epsilon} - 1)p(H|S') + \delta\Big] \end{aligned}$$

Both results imply that:

$$\big|p(H|S) - p(H|S')\big| \leq (e^{\epsilon} - 1)p(H|S') + \delta \qquad (14)$$

Writing:

$$\begin{aligned} \mathcal{J}(Z_{trn};\,H) &= ||p(Z_{trn},H)\,,\;\; p(Z_{trn}) \cdot p(H)||_{\mathcal{T}} \\ &= \mathbb{E}_{Z_{trn}}\; ||p(H|Z_{trn})\,,\;\; p(H)||_{\mathcal{T}} \\ &= \frac{1}{2}\mathbb{E}_{Z_{trn}}\Big|\Big|\mathbb{E}_{Z'_{trn}}\big[p(H|Z_{trn}) - p(H|Z'_{trn})\big]\Big|\Big|_1 \\ &\leq \frac{1}{2}\mathbb{E}_{Z_{trn},Z'_{trn}}\Big|\Big|p(H|Z_{trn}) - p(H|Z'_{trn})\Big|\Big|_1 \end{aligned}$$

The last inequality follows by convexity. Next, let S_{m-1} be a sample that contains $m-1$ observations drawing i.i.d. from $p(z)$. Then:

$$\begin{aligned} \mathcal{J}(Z_{trn};\,H) &\leq \frac{1}{2}\mathbb{E}_{Z_{trn},Z'_{trn}}\Big|\Big|p(H|Z_{trn}) - p(H|Z'_{trn})\Big|\Big|_1 \\ &= \frac{1}{2}\mathbb{E}_{Z_{trn},Z'_{trn}}\Big|\Big|\mathbb{E}_{S_{m-1}}\big[p(H|Z_{trn},S_{m-1}) - p(H|Z'_{trn},S_{m-1})\big]\Big|\Big|_1 \\ &\leq \frac{1}{2}\mathbb{E}_{S,S'}\Big|\Big|p(H|S) - p(H|S')\Big|\Big|_1, \end{aligned}$$

where S, S' are two adjacent samples.

Next, we expand the ℓ_1 distance and use Eq 14:

$$\begin{aligned} \mathcal{J}(Z_{trn};\,H) &\leq \frac{1}{2}\mathbb{E}_{S,S'}\Big|\Big|p(H|S) - p(H|S')\Big|\Big|_1 \\ &= \frac{1}{2}\mathbb{E}_{S,S'}\sum_{H\in\mathcal{H}}\big|p(H|S) - p(H|S')\big| \\ &\leq \frac{1}{2}\mathbb{E}_{S,S'}\sum_{H\in\mathcal{H}}\big[(e^{\epsilon} - 1)p(H|S') + \delta\big] \\ &= \frac{e^{\epsilon} - 1 + \delta}{2} \end{aligned}$$

Finally, the desired bound follows by combining the last inequality with Theorem 4.

K Proof of Corollary 2

It has been shown in Alabdulmohsin (2015) that the supremum generalization risk is achieved (arbitrarily well) using the following binary-valued loss:

$$L^{\star}(z;H) = \mathbb{I}\big\{p(Z_{trn} = z|H) \geq p(Z_{trn} = z)\big\} \qquad (15)$$

Therefore, if an algorithm is (ϵ, δ) robustly generalizing, let the adversary $\mathcal{A}$ (or equivalently the parametric loss $L(\cdot; H)$) be fixed to the one given by Eq. 15. Hence, we have by definition of robust generalization:

$$p\Big\{\big|\mathbb{E}_{Z\sim p(z)}L^\star(Z;H) - \frac{1}{m}\sum_{Z_i\in S} L^\star(Z_i;H)\big| \leq \epsilon\Big\} \geq 1-\delta, \tag{16}$$

Therefore:

$$\Big|\mathbb{E}_{S,H}\big[\mathbb{E}_{Z\sim p(z)}L^\star(Z;H) - \frac{1}{m}\sum_{Z_i\in S} L^\star(Z_i;H)\big]\Big| \leq \epsilon + \delta$$

Because $L^\star(\cdot; H)$ achieves the maximum possible generalization risk in expectation (Alabdulmohsin, 2015), we have the uniform generalization bound $\mathcal{J}(Z_{trn};\, H) \leq \epsilon+\delta$. Hence, (ϵ, δ) robust generalization implies a uniform generalization at the rate $\epsilon + \delta$.

The proof of the converse follows from our concentration bound in Theorem 4, which shows that uniform generalization in expectation implies a generalization in probability. In particular, any algorithm that generalizes uniformly with rate τ is (ϵ, γ) robustly generalizing, with $\gamma = (5/2)(\tau + \sqrt{\log 9/(25m)})/\epsilon$.

L Proof of Theorem 5

Before we prove the statement of the theorem, we begin with the following lemma:

Lemma 1. *Let the observation space $\mathcal{Z}$ be the interval $[0, 1]$, where $p(z)$ is continuous in $[0, 1]$. Let $H \subseteq S_m : |H| = k$ be a set of k examples picked at random without replacement from the sample S_m. Then $\mathcal{J}(Z_{trn};\, H) = \frac{k}{m}$.*

Proof. First, we note that $p(Z_{trn}|H)$ is a *mixture* of two distributions: one that is uniform in H with probability k/m, and the original distribution $p(z)$ with probability $1 - k/m$. By Jensen's inequality, we have $\mathcal{J}(Z_{trn};\, H) \leq k/m$. Second, let the parametric loss be $L(\cdot; H) = \mathbb{I}\{Z \in H\}$. Then, $|R_{gen}(\mathcal{L})| = \frac{k}{m}$. By Theorem 1, we have $\mathcal{J}(Z_{trn};\, H) \geq |R_{gen}(\mathcal{L})| = k/m$. Both bounds imply the statement of the lemma. □

Now, we prove Theorem 5. Consider the example where $\mathcal{Z} = [0, 1]$ and suppose that the observations $Z \in \mathcal{Z}$ have a continuous marginal distribution. Because t is a rational number, let the sample size m be chosen such that $k = t\, m$ is an integer.

Let $\{Z_1, \ldots, Z_m\}$ be the training set, and let the hypothesis H be given by $H = \{Z_1, \ldots, Z_k\}$ with some probability $\delta > 0$ and $H = \{\}$ otherwise. Here, the k instances $Z_i \in H$ are picked uniformly at random without replacement from the sample S_m. To determine the variational information between Z_{trn} and H, we consider the two cases:

1. If $H \neq \{\}$, then $||p(Z_{trn})\,,\ p(Z_{trn}|H)||_{\mathcal{T}} = t$ as proved in Lemma 1. This happens with probability δ by construction.
2. If $H = \{\}$ then $p(Z_{trn}|H) = p(Z_{trn})$. Hence, we have $||p(Z_{trn})\,,\ p(Z_{trn}|H = \{\})||_{\mathcal{T}} = 0$. This happens with probability $1 - \delta$.

So, by combining the two cases above, we deduce that:

$$\mathcal{J}(Z_{trn};\, H) = \mathbb{E}_H\ ||p(Z_{trn})\,,\ p(Z_{trn} \mid H)||_{\mathcal{T}} = t\ \delta.$$

Therefore, $\mathcal{L}$ generalizes uniformly with the rate $t\delta$. Next, let the loss $L(\cdot; H)$ be given by $L(Z\ ;\ H) = \mathbb{I}\big\{Z \in H\big\}$. With this loss:

$$p\Big\{\big|R_{emp}(H; S_m) - R_{true}(H)\big| = t\Big\} = \delta = \frac{\mathcal{J}(Z_{trn};\, H)}{t},$$

which is the statement of the theorem.

M Proof of Proposition 4

This proposition is proved in Corollary 1.

References

Alabdulmohsin, I. (2015). Algorithmic stability and uniform generalization. In *NIPS*, pages 19–27.

Audibert, J.-Y. and Bousquet, O. (2007). Combining PAC-Bayesian and generic chaining bounds. *JMLR*, 8:863–889.

Bartlett, P. L. and Mendelson, S. (2002). Rademacher and gaussian complexities: Risk bounds and structural results. *JMLR*, 3:463–482.

Bassily, R. and Freund, Y. (2016). Typicality based stability and privacy. arXiv:1604.03336 [cs.LG].

Bassily, R., Nissim, K., Smith, A., Steinke, T., Stemmer, U., and Ullman, J. (2016). Algorithmic stability for adaptive data analysis. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC)*, pages 1046–1059.

Blumer, A., Ehrenfeucht, A., Haussler, D., and Warmuth, M. K. (1989). Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM (JACM)*, 36(4):929–965.

Boucheron, S., Lugosi, G., and Bousquet, O. (2004). Concentration inequalities. In *Advanced Lectures on Machine Learning*, pages 208–240. Springer.

Bousquet, O., Boucheron, S., and Lugosi, G. (2004). Introduction to statistical learning theory. In Bousquet, O., von Luxburg, U., and Rätsch, G., editors, *Advanced Lectures on Machine Learning*, volume 3176, pages 169–207.

Bousquet, O. and Elisseeff, A. (2002). Stability and generalization. *JMLR*, 2:499–526.

Cover, T. M. and Thomas, J. A. (1991). *Elements of information theory*. Wiley & Sons.

Csiszár, I. (1972). A class of measures of informativity of observation channels. *Periodica Mathematica Hungarica*, 2:191–213.

Csiszár, I. (2008). Axiomatic characterizations of information measures. *Entropy*, 10:261–273.

Cummings, R., Ligett, K., Nissim, K., Roth, A., and Wu, Z. S. (2016). Adaptive learning with robust generalization guarantees. In *Proceedings of the 29th Annual Conference on Learning Theory (COLT)*.

David, O., Moran, S., and Yehudayoff, A. (2016). Supervised learning through the lens of compression. In *NIPS*.

Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., and Roth, A. (2015). Preserving statistical validity in adaptive data analysis. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC)*, pages 117–126.

Dwork, C. and Roth, A. (2013). The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4):211–407.

Hardt, M., Recht, B., and Singer, Y. (2016). Train faster, generalize better: Stability of stochastic gradient descent. In *Proceedings of the 33rd international conference on Machine learning (ICML)*.

Luntz, A. and Brailovsky, V. (1969). On estimation of characters obtained in statistical procedure of recognition. *Technicheskaya Kibernetica*, 3(6).

McAllester, D. (2003). PAC-Bayesian stochastic model selection. *Machine Learning*, 51:5–21.

Reid, M. D. and Williamson, R. C. (2009). Generalised Pinsker inequalities. In *COLT*.

Russo, D. and Zou, J. (2016). Controlling bias in adaptive data analysis using information theory. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*.

Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.

Shalev-Shwartz, S., Shamir, O., Srebro, N., and Sridharan, K. (2010). Learnability, stability and uniform convergence. *JMLR*, 11:2635–2670.

Stigler, S. M. (1986). *The history of statistics: The measurement of uncertainty before 1900*. Harvard University Press.

Vapnik, V. and Chapelle, O. (2000). Bounds on error expectation for support vector machines. *Neural Computation*, 12(9):2013–2036.

Vapnik, V. N. (1999). An overview of statistical learning theory. *Neural Networks, IEEE Transactions on*, 10(5):988–999.